MARK D. FOWLER, Bar No. 124235
mark.fowler@dlapiper.com
CLAYTON THOMPSON (Admitted Pro Hac Vice)
clayton.thompson@dlapiper.com
DAVID ALBERTI, Bar No. 220625
david.alberti@dlapiper.com
CHRISTINE K. CORBETT, Bar No. 209128
christine.corbett@dlapiper.com
YAKOV M. ZOLOTOREV, Bar No. 224260
yakov.zolotorev@dlapiper.com
CARRIE L. WILLIAMSON, Bar No. 230873
carrie.williamson@dlapiper.com

DLA PIPER US LLP
2000 University Avenue
East Palo Alto, CA 94303-2214
Tel: 650.833.2000
Fax: 650.833.2001

Attorneys for Defendant and Counterclaim Plaintiff
SUN MICROSYSTEMS, INC.

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

SAN FRANCISCO DIVISION

NETWORK APPLIANCE, INC.,

Plaintiff-Counterclaim Defendant,

v.

SUN MICROSYSTEMS, INC.,

Defendant-Counterclaim Plaintiff.

CASE NO. C-07-06053-EDL

SUN MICROSYSTEMS, INC.'S RESPONSIVE CLAIM CONSTRUCTION BRIEF

Claim Construction Hearing:
Date: August 27, 2008
Time: 9:00 a.m.
Judge: Hon. Elizabeth D. Laporte

TABLE OF CONTENTS

TABLE OF CONTENTS
(continued)

Page

TABLE OF AUTHORITIES

Page

CASES

TABLE OF AUTHORITIES
(continued)

Page

I. INTRODUCTION

This brief addresses Sun's United States Patent Nos. 5,925,106 ("the '106 patent"), 5,459,857 ("the '857 patent"), 5,749,095 ("the '095 patent"), and 6,873,630 ("the '630 patent"). More specifically, this brief addresses the claim terms in these four patents that were identified for construction in the parties' June 16, 2008, Joint Report Regarding 20 Claim Terms For Claim Construction (the "Report"). Although, pursuant to the Court's June 5, 2008, Case Management Order, and as stated at page 1 of the Report, five Sun patents were to be construed by the Court at this time, the Report does not identify any claim term as requiring construction in the fifth Sun patent, United States Patent No. 6,681,261.

II. U.S. PATENT NO. 5,925,106

A. Technology Background.

The '106 patent concerns obtaining and displaying information about a server on a network. '106 patent, col. 1:1-3, col. 2:23-33. The basic problem addressed by the '106 patent concerns a computer user who is accessing a network consisting of multiple servers. Specifically, the user may need to obtain and display information about the particular server from which the user is, at that point in time, accessing information. *Id.*, Abstract, col. 1:53-2:20; Declaration of Dr. Martin Kaliski ("Kaliski Decl."), ¶ 7.

Servers can be accessed over the Internet ***or*** over other local computer networks using domain names. '106 patent, col. 1:54-67; Kaliski Decl., ¶ 7. Domain names are names that have numerical IP addresses associated with them. *Id.*

Servers can be distinguished from other servers on a network by "server identification data" - data that more descriptively differentiates one server from another. '106 patent, Abstract, col. 2:9-40, col. 5:33-67, Figs. 3B, 3C, 4A, 4B. Server identification data includes information describing or identifying the server. *Id.*

The specification describes a preferred embodiment of the invention in the context of a World Wide Web ("WWW") server. '106 patent, col. 1:34-37, col. 4:24-30. However, the specification expressly states the WWW context is just one example, as the invention covers obtaining information from any type of network server, such as a server on a private network. *Id.*,

col. 1:32-34, col. 2:23-56, col. 4:5-10, 24-26; Kaliski Decl., ¶ 8. The patent teaches how to request descriptive information about a server by, for example, the client (user) computer querying the server for such information. '106 patent, col. 6:22-26; Kaliski Decl., ¶ 8. The patent teaches how to accomplish such queries by the client to the server. '106 patent, col. 6:22-36; Kaliski Decl., ¶ 8. The patent discloses displaying information about the server using software resident on both the client and the server. '106 patent, Figs. 5, 6, col. 6:22-col. 7:55; Kaliski Decl., ¶ 8.

Some of the patent's dependent claims are directed to the WWW, uniform resource locators ("URLs"), hypertext transfer protocol ("HTTP") and specific query commands (*e.g.*, claims 5-8, 11-13, 18-21, 24-26, 31-34, 37-39). However, other claims, including all of the patent's independent claims, relate more generally to providing information about a server over a network to a client computer and then displaying the information on the client computer (*e.g.*, claims 1, 4, 9, 10, 14, 17, 22, 23, 27, 30, 35, 36).

B. Representative Claim.

Claim 1 of the '106 patent is reproduced below with the disputed terms identified in bold.

> 1. An information access apparatus configured to access information; said information access apparatus having a central processing unit (CPU), a memory, a network interface to provide access to a network and a display device; said information access apparatus utilizing a ***domain name*** to access said information existent on said network; said information access apparatus comprising:
>
> > an access mechanism configured to obtain ***server identification data*** existent on said network and accessible by use of said ***domain name***; said ***server identification data*** including descriptive information about a server; and
> >
> > a display mechanism configured to display said ***server identification data*** on said display device.

'106 patent, col. 9:45-57 (emphasis added).

C. Disputed Claim Terms To Be Construed.

1. "domain name"

The term "domain name" appears in each independent claim of the '106 patent.

Sun's proposed construction	NetApp's proposed construction
A name that has a numerical IP address associated with it.	A third-party approved name of a website on the Internet, i.e., a registered domain name.

The critical issue here is whether the term "domain name" is limited to a single type of domain name – *i.e.*, the registered domain name of an Internet website (NetApp's construction) – or whether the term also includes other types of known domain names (Sun's construction).

a. Sun's Construction Is Consistent With The Intrinsic Evidence.

The only construction consistent with the intrinsic evidence is Sun's construction. NetApp's narrow construction is not supported (and certainly is not required) by the intrinsic record. Simply put, whether a domain name has been registered or "approved" by "a third-party," or identifies an Internet website, is irrelevant to practicing the claimed invention.

(1) The Specification Supports Sun's Construction.

The specification repeatedly explains that use of the Internet, the WWW, URLs and HTTP is merely representative of the technology covered by the invention, and that such use merely constitutes a preferred embodiment of the patent. '106 patent, col. 1:32-37 ("Although the invention covers information access and information provider apparatus, WWW browser and WWW server applications are representative of the technology. As such, the majority of this application describes the invention within the context of a preferred embodiment utilizing WWW browser and WWW server applications"), col. 4:24-29 ("Although the invention covers information access and information provider apparatus, methods and computer program products, WWW browser and WWW server apparatus and applications are representative of the technology. As such, a preferred embodiment utilizing WWW browser and WWW server applications are described"). Thus, contrary to NetApp's assertion, the specification dictates that the invention cannot be limited to the context of the Internet and the WWW.

The breadth of the invention is reflected in its description in the Summary of the Invention section of the specification. The Summary of the Invention does not describe the invention in relation to the Internet, the WWW or any other specific technology. *Id.*, col. 2:24-59. Instead, the invention is described as applying to any type of network, any type of server and any type of

client computer. *Id.* This description encompasses both the Internet and WWW, as well as private networks using domain names not registered by a third party. Kaliski Decl., ¶ 13. In this regard, it is common for private networks to use domain names that are not "third-party approved" or registered. *Id.*

Adopting NetApp's proposed construction – which would ignore both the express statement in the specification that the invention is not limited to the WWW, and the broad language of the Summary of the Invention – would violate the rule that "care must be taken to avoid reading material appearing in the specification into the claims." *Interactive Gift Express, Inc. v. Compuserve, Inc.*, 256 F.3d 1323, 1331 (Fed. Cir. 2001); *see also Phillips v. AWH Corp.*, 415 F.3d 1303, 1323 (Fed. Cir. 2005) ("we have expressly rejected the contention that if a patent describes only a single embodiment, the claims of the patent must be construed as being limited to that embodiment"); *Decisioning.com, Inc. v. Federated Dept. Stores, Inc.*, 527 F.3d 1300, 1309 (Fed. Cir. 2008) (refusing to limit claims to preferred embodiment); *Saunders Group, Inc. v. Comfortrac, Inc.* 492 F.3d 1326, 1333 (Fed. Cir. 2007) (refusing to limit claims to sole embodiment described in specification).

NetApp quotes *Microsoft Corp. v. Multi-Tech Sys., Inc.*, 357 F.3d 1340, 1347 (Fed. Cir. 2004). However, the passage quoted by NetApp refers to a case "[w]hen the specification makes clear that an invention does not include a particular feature . . ." *Id.* Here, by contrast, not only does the specification not exclude Sun's construction, it expressly states the invention is not limited to NetApp's construction.

Nothing about the operation of the invention requires using "third-party approved" or "registered" domain names. Kaliski Decl., ¶ 15. For instance, the '106 patent explains server identification data can be retrieved using a conventional database that associates domain names with server identification data. '106 patent, col. 7:56-8:4; Kaliski Decl., ¶ 15. Such a database need not be maintained by a third party or use only registered domain names. Kaliski Decl., ¶ 15. Rather, the invention can be used on ***any*** type of network and by ***any*** type of client to retrieve information about a server accessed through ***any*** type of domain name. *Id.*

NetApp cites parts of the specification in purported support for its narrow construction.

However, each citation blindly ignores that the parts of the specification cited by NetApp are merely discussing a non-exclusive preferred embodiment involving the WWW. Properly viewed in that context, it is not surprising that the language cited by NetApp discusses domain names in the context of the WWW. As established above, however, the specification itself establishes that the invention is not limited to the context of domain names in the WWW.

NetApp also points to certain documents pertaining to the WWW that are incorporated by reference into the specification. '106 patent, col. 1:25-44; Walter Decl., Exs. J, K, L, M. However, in citing these documents, NetApp again blindly ignores that the specification states, in the same section these WWW-related documents are cited, that WWW technology is merely representative of a preferred embodiment of the broader invention. *Id.*

(2) The Claim Language Supports Sun's Construction.

Notably, NetApp's brief entirely ignores the language of the claims. This is perhaps explained by the fact that the plain language of the claims mandates the rejection of NetApp's proposed construction.

All the independent claims – and even some of the dependent claims – recite systems and methods for providing server information to the client over a network – any network. *See, e.g.,* '106 patent, claim 1 ("An information access apparatus configured to access information . . . having a network interface to provide access to a network"), 2-4, 9 ("A server apparatus having . . . access to a network"), 10, 14 ("on a network"), 17, 22 ("over a network"), 23, 27 ("access to a network"), 30, 35 ("access to a network") and 36. In contrast, some of the dependent claims add the use of the WWW, URLs and HTTP. *See, e.g.,* '106 patent, claim 5 ("The information access apparatus of claim 4 wherein said access mechanism further comprises a world wide web (WWW) browser apparatus configured to use a hypertext transfer protocol (HTTP) to effect said query mechanism and said receiver mechanism, said HTTP utilizing a uniform resource locator (URL)"), 6-8, 11 (requires usage of WWW, HTTP and URLs), 12-13, 18 (requires usage of HTTP and URLs) 19-21, 24 (same), 25-26, 31 (requires usage of WWW, HTTP and URLs), 32-34, 37 (same) and 38-39.

NetApp's proposed construction, which limits domain names to registered Internet (*i.e.*,

WWW) domain names, thus impermissibly renders much of the language in the dependent claims superfluous. Specifically, by limiting the claim term "domain name" to "registered Internet domain names," NetApp's construction would impose WWW, URL and HTTP requirements in ***all*** of the independent claims, thereby rendering meaningless the addition of the WWW, URL and HTTP claim limitations in the narrow dependent claims. Kaliski Decl., ¶ 14.

Such a requirement would violate the principle of claim differentiation, which provides that where some claims are broad and others narrow, the narrow claim limitations cannot be read into the broader claims. *Uniroyal Inc. v. Rudkin-Wiley Corp.*, 837 F.2d 1044, 1054-55 (Fed. Cir. 1988) (District Court erred in reading claims 1 and 2 as if they contained the same limitations when only claim 1 contained the narrow limitation). Because many of the dependent claims expressly recite WWW, URL and HTTP usage, Sun's construction is preferred because all claim terms are presumed to have independent meaning and a construction that gives independent meaning to all terms is preferred over one that does not. Thus, constructing a claim term in a way that renders other claim terms superfluous – as NetApp's construction does – is incorrect. *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1119 (Fed. Cir. 2004); *Merck & Co. v. Teva Pharm. USA, Inc.*, 395 F.3d 1364, 1372 (Fed. Cir. 2005) (same); *Phillips*, 415 F.3d at 1324 ("[t]he inclusion of ... a specific limitation on the term 'baffles' in claim 2 makes it likely that the patentee did not contemplate that the term 'baffles' [in claim 1] already contained that limitation.").

(3) The Prosecution History Supports Sun's Construction.

During prosecution of the '106 patent, the applicant expressly defined the term "domain name" in a manner consistent with Sun's construction and inconsistent with NetApp's construction. Specifically, in a December 22, 1997, Amendment, the applicant explained that a "patentee can be his own lexicographer" and expressly defined three terms: "server identification information," "domain name" and "descriptive information about a server." The applicant defined "domain name" as follows:

> The "Domain Name" term is used in accordance with standard usage in the field. It is provided to a nameserver to return an IP address.

Walter Decl., Ex. E, 12/22/97 Amendment, at 5-6.

A "nameserver" simply is a mechanism for associating names with IP addresses. *Id.*, at 4-5. Therefore, consistent with the applicant's express definition, a "domain name" is a name with an IP address associated with it. Nowhere in the applicant's definition of "domain name" is there any reference to, much less a requirement of, "third-party approval" or "registration."

Furthermore, contrary to NetApp's assertion, the applicant never argued a domain name had to be "registered" in order to distinguish prior art. In the December 22, 1997, Amendment, the applicant distinguished the claimed invention from U.S. Patent No. 5,664,185 ("Landfield") based on the fact that Landfield did not return or display server identification data. Walter Decl., Ex. E, at 6-9. The issue was not whether Landfield disclosed "domain names" or "registered domain names." In fact, the applicant acknowledged Landfield disclosed domain names when he explained domain names are "provided to a nameserver to return an IP address." *Id.*, at 6. In any event, there certainly is nothing clear or unmistakable in the prosecution history that limits the term "domain name" to "registered domain names." *Omega Eng'g, Inc., v. Raytek Corp.*, 334 F.3d 1314, 1325-26 (Fed. Cir. 2003) (citation omitted) ("for prosecution disclaimer to attach, our precedent requires that the alleged disavowing actions or statements made during prosecution be both clear and unmistakable").

b. The Extrinsic Evidence Supports Sun's Construction.

None of the extrinsic evidence cited by NetApp supports NetApp's narrow construction. Particularly, none of the dictionary definitions cited mention – much less require – third party approval or registration for a name to be considered a "domain name." Rather, consistent with Sun's definition, the Microsoft Press Computer Dictionary, 3d Ed. defines "domain name" to be "[a]n address of a network connection that identifies the owner of that address in a hierarchical format: *sever.organization.type* …" Walter Decl., Ex. F, p. 158; *see also* Webster's New World Dictionary Of Computer Terms, 8th ed. (2000) (defining "domain" as "[i]n a computer network, a group of computers that are administered as a unit…."; "domain name" as "[i]n the system of domain names used to identify individual Internet computers, a single word or abbreviation that makes up part of a computer's unique name (such as watt.seas.virginia.edu). …") Walter Decl.,

Ex. G. Thus, while NetApp argues "standard usage" of the term "domain name" connotes only registered domain names, NetApp fails to cite a single dictionary definition that includes the word "register," "registration," "third party approval" or any comparable term.

Even the RFC reference cited by NetApp as authoritative on domain names explains that not all domain names need to be DNS-type or registered domain names:

> **The terms "domain" or "domain name" are used in many contexts beyond the DNS described here.** Very often, the term domain name is used to refer to a name with structure indicated by dots, but no relation to the DNS. This is particularly true in mail addressing [Quarterman 86].

Walter Decl., Ex. K, Section 2.1, "The history of domain names" (emphasis added). Thus, "domain names" are used in many contexts other than "registered" DNS-type domain names.

Indeed, NetApp itself regularly uses the term "domain name" to describe "a name that has a numerical IP address associated with it" and is ***not*** registered or third-party approved. For example, in documentation describing the accused products (*e.g.*, Data ONTAP 7.2), NetApp's own network administration guide explains that "domain names" can include things other than DNS domain names, such as NIS domain names, which need not be registered or third-party approved. Declaration of Carrie Williamson In Support Of Sun's Responsive Claim Construction Brief, Ex. A, Data ONTAP® 7.2, Network Management Guide (March 2007); Kaliski Decl., ¶ 16. In fact, NetApp repeatedly uses the phrase "NIS domain name" throughout the Network Management Guide. *Id.*, at 101, 103, 104, 105, 107. NetApp's use of the term "domain name" in a manner consistent with Sun's construction and contrary to its own construction supports adopting Sun's construction. *Serio-US Indus., Inc. v. Plastic Recovery Tech. Corp.*, 459 F.3d 1311, 1319 (Fed. Cir. 2006) ("a trial court may consult the accused device for context that informs the claim construction process"); *Lava Trading, Inc. v. Sonic Trading Mgmt., LLC*, 445 F.3d 1348, 1350 (Fed. Cir. 2006) ("knowledge of [an accused] product or process provides meaningful context for the first step of the infringement analysis, claim construction").

NetApp's reliance on the "Telnet" art also is misplaced. As an initial matter, Dr. Almeroth's discussion of "Telnet" is superficial and certainly does not support a conclusion that all limitations of any claim are anticipated. Kaliski Decl., ¶ 17. Indeed, Dr. Almeroth's invalidity

"analysis" is limited to the following single conclusory sentence: "If Sun's construction were correct, the use of Telnet with the '/etc/hosts' file would incorporate every limitation of Claim 1." Almeroth Decl., at 9. In any event, as a matter of law, NetApp's invalidity contention cannot impact a claim construction clearly required by the intrinsic evidence. *Phillips*, 415 F.3d at 1327.

c. The Federal Circuit Cases Cited By NetApp As Extrinsic Evidence Are Inapposite.

NetApp cites several Federal Circuit cases as extrinsic evidence in support of its proposed construction. *In re Oppedahl & Larson LLP*, 373 F.3d 1171 (Fed. Cir. 2004) (affirming the USPTO's refusal to register "patents.com" as a trademark); *Resonate Inc. v. Alteon Websystems, Inc.*, 338 F.3d 1360 (Fed. Cir. 2003) (patent infringement action regarding U.S. Patent No. 5,774,660). The cases, however, do not support limiting the term "domain name" in the '106 patent to third-party-approved registered Internet addresses. First, the application for the '106 patent was filed in April 1996, over seven years before either of these cases was decided. Thus, neither case is extrinsic evidence of what one of ordinary skill in the art understood at the time of the patent. Second, *Oppedahl* and *Resonate* concerned, respectively, a trademark relating to a web site address and a patent covering Internet server technology. Nowhere in either case does the Federal Circuit state or imply the term "domain name" *exclusively* refers to registered Internet addresses. At most, the Federal Circuit implicitly acknowledged that registered Internet addresses are understood as falling within the scope of the broader term "domain name." NetApp's citation to the dissent in the order denying a rehearing *en banc* in *Teva Pharm.* also is inapposite because its passing discussion of cybersquatting does not support the proposition that the term "domain name" is limited to registered Internet addresses. *Teva Pharm. USA, Inc. v. Pfizer Inc.*, 405 F.3d 990, 997 (Fed. Cir. 2005).

Similarly, NetApp's citation to federal trademark law as extrinsic evidence does not support limiting the term "domain name" in the '106 patent to third-party approved Internet addresses. As an initial matter, the statutory trademark definition of "domain name" cited by NetApp was not enacted until November 1999, over three years after the application for the '106 patent was filed. As such, it is not extrinsic evidence of the understanding of one of skill in the

art at the time of filing. Indeed, while the statutory definition of "domain name" in trademark law might be known to lawyers who practice in that field, NetApp fails to show that one of ordinary skill in the art would know of this legislation, much less that this legislation somehow reflects the understanding of one of skill in the art. Moreover, Congress narrowly defined "domain name" in the Lanham Act because it was not seeking to extend trademark liability to internal domain names and private IP addresses.

2. "server identification data"

This phrase appears in each independent claim of the '106 patent.

Sun's proposed construction	NetApp's proposed construction
Information that uniquely identifies one server from other servers and can be seen by a user.	Human-friendly information identifying a specific web server designed not to be intimidating to inexperienced users of the World Wide Web.

The distinctions between the parties' respective constructions are whether the term "server identification data": (1) is limited to information about web servers on the World Wide Web (NetApp's construction) or, instead, also encompasses information about other types of servers (Sun's construction); and (2) is limited to "human friendly" information and that is "designed not to be intimidating" (NetApp's construction).

a. The Specification Supports Sun's Construction.

The Summary of the Invention section of the specification describes the invention as a way of "providing server-specific information to a computer user." '106 patent, col. 2:24-28. The server-specific information, called "sever identification data," is obtained using a domain name and can be displayed to a user using a display device. *Id.*, col. 2:28-31. The specification describes "server identification data" as including "descriptive information about the server." *Id.*, col. 2:31-33. Sun's proposed construction captures all of these aspects of what the specification describes as "the invention."

NetApp's construction conflicts with the specification in numerous ways. First, it limits this term to websites on the WWW. As established above in connection with the claim term "domain name," the specification specifically states that the patent is not limited to the WWW.

Second, rather than focusing on what the '106 patent describes as "the invention," NetApp selectively culls excerpts related to the "preferred embodiment" and cobbles together an ambiguous construction requiring "human friendly" information that is "designed to not be intimidating to inexperienced users of the World Wide Web." There is no intrinsic evidence that these phrases were intended by the applicant to constitute a limitation on the invention and the claims. As such, NetApp falls far short of establishing that these statements can be read into the claims as limitations. *Interactive Gift Express, Inc.*, 256 F.3d at 1331; *Teleflex, Inc. v. Ficosa North Am. Corp.*, 299 F.3d 1313, 1325-26 (Fed. Cir. 2002) (limitations cannot be read into the claims from the specification absent an express intent to do so).

NetApp quotes three sentences from the specification referring to the "invention." NetApp Op. Br., at 12. Notably, none of the quoted passages refers either to the WWW or to being "designed not to be intimidating to inexperienced users of" the WWW (or any other words to the effect). Thus, NetApp's reliance on *Verizon Serv. Corp. v. Vonage Holdings Corp.*, 503 F.3d 1295, 1308 (Fed. Cir. 2007) is entirely misplaced, as these parts of NetApp's proposed construction are not described in the specification as constituting "the invention."

Moreover, NetApp's proposed construction does not serve the purposes of claim construction. Rather than providing guidance and meaning to the jury, NetApp's proposed construction only raises subjective questions, such as what constitutes "human friendly" and what constitutes something "designed not to be intimidating"?

b. The Claims Support Sun's Construction.

As established above in connection with the claim term "domain name," NetApp's proposed restriction of this claim term to websites on the WWW is incorrect because: (1) no such limitation exists in the independent claims; and (2) if adopted, NetApp's construction would improperly render the WWW limitations in the dependent claim superfluous. Sun's proposed construction, in contrast, is supported by existing claim language: (1) specifying that "server identification data" "include[e] descriptive information about a server"; and (2) reciting that the "server identification data" can be "displayed" on a "display device" (the only purpose of which would be for a "user" to view the data). '106 patent, claims 1, 14. Accordingly, while Sun's

construction is supported by the claim language, NetApp's construction is unsupported by, and is in fact at odds with, the claim language.

c. The Prosecution History Supports Sun's Construction.

During prosecution of the '106 patent, the applicant expressly defined the term "server identification information" in a manner consistent with Sun's construction and inconsistent with NetApp's construction. In the December 22, 1997, Amendment, the applicant expressly defined "server identification information" as follows:

> The "Server Identification Information" term is used throughout the application to mean descriptive information about a server that can be seen by the user - in particular text page 9 lines 3-4, 9-11, 17-21,27-28; page 9 line 33 - page 10 line 2; page 11 lines 8-15). This descriptive information is not an IP address that is returned by a nameserver.

Walter Decl., Ex. E, at 5-6. This definition of "server identification information" is wholly consistent with Sun's construction. By contrast, nowhere in the applicant's definition is there any reference to, or requirement of, a WWW server, or that the information be "human-friendly" or "designed not to be intimidating."

NetApp's own reliance on the prosecution history is both flawed and misplaced. First, in the section of the file history cited by NetApp, the applicant stated the claimed "descriptive information" "is not an IP address that is returned by a nameserver." Walter Decl., Ex. E, at 5-6. However, the term being construed here is not the claimed "descriptive information about a server"; rather, it is the claimed "server identification data." In this regard, all of the independent claims recite that "server identification data" *includes* "descriptive information about a server." '106 patent, claims 1, 9, 14, 22, 27, 35. Therefore, the term "server identification data" is inherently broader than the term "descriptive information about a server." Thus, any alleged limitations on the scope of "descriptive information about a server" would not necessarily apply to the broader term "server identification data."

Second, and perhaps more important, Sun does not contend that an "IP address" is either the "server identification data" or the "descriptive information about a server." Sun has no objection to the Court stating in its claim construction order that an IP address does not itself constitute either limitation. This moots NetApp's main argument.

Third, contrary to NetApp's assertion, the applicant did not disclaim the domain name, hypertext links, URLs, web page titles and bookmarks as being part of the "descriptive information." Rather, the applicant stated that such information would be presented along with other descriptive information about the server. Walter Decl., Ex. E, at 4. In any event, in order to further moot NetApp's argument, Sun does not object to the Court stating in its claim construction order that the domain name, hypertext links, URLs, web page titles and bookmarks do not by themselves satisfy the "server identification data" or "descriptive information about a server" limitations. This entirely moots NetApp's prosecution history argument.

III. U.S. PATENT NO. 5,459,857

A. Technology Background.

The '857 patent is directed to a fault tolerant disk array system in which the probability of failure is greatly minimized. '857 patent, col. 1:8-10. The '857 patent accomplishes this through multiple forms of redundancy. The data storage system illustrated in Figure 1 of the patent is duplicated below.

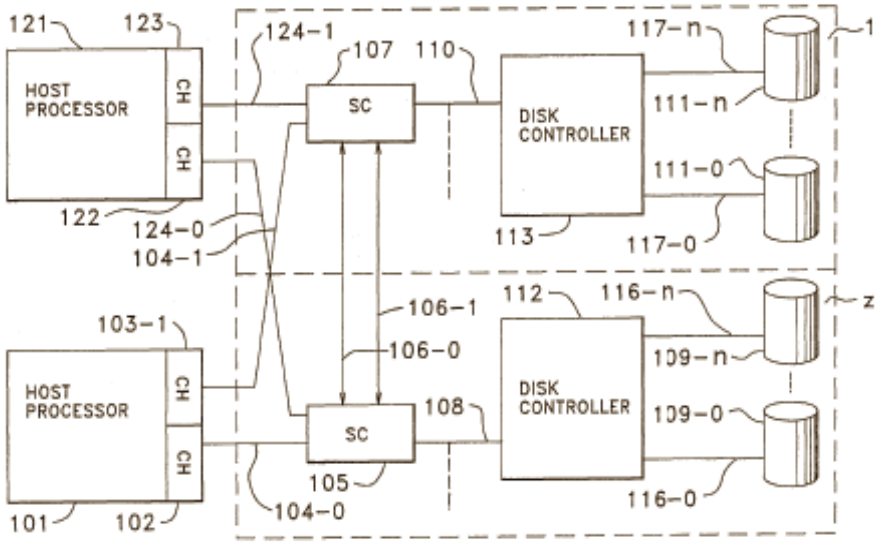

FIG.1.

'857 patent, Fig. 1, col. 4:26-54. As illustrated above, the storage system has two data storage subsystems, 1 and 2. *Id.* In the preferred embodiment, the storage controller(s) 105 and 107, disk controller(s) 112 and 113, and disk drives 109 and 111 are duplicates, with the same information

stored in each duplicate. *Id.*

Information is stored in a redundant manner on each data storage system using redundant arrays of disk drives (*e.g.,* "RAID" arrays), which are configured into "redundancy groups" of "N+M" disk drives (where "N" is a number of (N) disk drives, and "M" is a number of (M) disk drives). '857 patent, col. 3:1-5. As in typical RAID systems, the N+M disk drives are used to store one N data segment on each of N physical tracks, and to store M redundancy segments, one on each M physical tract per logical track in the group. *Id.*, col. 3:5-14. The redundancy segments can be used to reconstruct data from one or more failed disks. By way of example, a failure of a number of the disk drives will not result in the loss of any information within that set, because the information can be reconstructed using the M redundancy segments. *Id.*, col. 3:25-38.

When a host processor (*e.g.,* host processor 101) transmits to one data storage system, the same data is stored in the duplicate system in a redundant manner. '857 patent, col. 4:39-54. This duplication and redundancy is accomplished without any additional effort by the host processor. Specifically, once the host processor writes to one data storage system, the two storage systems communicate with each other (without further effort from the host processor) to enable the information to be stored in both data storage systems. *Id.*, col. 11:21-12:13.

This design has the advantage that the failure of any single component (*e.g.,* a disk drive, disk controller, storage controller, channel) will not result in a system failure - *i.e.,* a situation in which information is lost or inaccessible. *See, e.g.,* Fig. 8, col. 5:1-19. Further, there will be no system failure even in the case of certain multiple component failures, including the complete destruction of one of the two duplicate storage systems, and including the simultaneous failure of individual disk drives in each of the duplicates. *Id.*

1. Writing Data Records.

The sequence of communications performed by the data storage subsystems in writing data records is important to the claim terms at issue. The '857 patent describes a single embodiment for writing duplicate data records through the data storage subsystems. The discussion of this embodiment, which follows the subheading "Data Write Operations" at column

11, line 9 of the specification, is explicit that the key stages between transmitting a data record and writing it to disk occur in the following sequences:

> 1. The host processor transmits the data record to a storage cluster control unit 501 within a storage control unit 105 of a first data storage subsystem. ('857 patent, col. 11:51-54.)
>
> 2. "Upon receipt of the data record," the storage cluster 501 stores the data record in cache memory 512. (*Id.*, col. 11:54-57.)
>
> 3. "Cache memory 512 writes the data record into volatile store and non-volatile store 520" and mapping information is updated. (*Id.*, col. 11:57-62.)
>
> 4. Storage cluster 501 transmits the written data record to storage control unit 107 of the second data storage subsystem. (*Id.*, col. 11:62-12:1.)
>
> 5. The storage control unit 107 stores the data record in its cache memory and updates its mapping table. (*Id.*, col. 11:62-12:1.)
>
> 6. The first storage control unit 105 writes the data record to disk (*i.e.*, to its redundancy group of disk drives). (*Id.*, col. 12:3-7.)
>
> 7. The second storage control unit 107 writes the data record to disk (*i.e.*, to its redundancy group of disk drives). (*Id.*, col. 12:7-9.)

This sequence of operations illustrates the following important aspects of the '857 patent invention: (i) the data record is transmitted from the first data storage subsystem to the second data storage subsystem ***after*** it has been written to ***cache memory***, but ***before*** it is written to a redundancy group ***of disks***; and (ii) the cache memory allows "back-end" operations of a system (*e.g.*, how and when data is moved from a storage control unit to disk) to be completely transparent to the host processor. '857 patent, col. 3:59-4:24, 6:64-7:1, 7:13-16. Because how and when data actually is stored on disk is irrelevant to the host processor, the data storage subsystems can notify the processor that a write operation is "complete" when a data record is stored in cache memory even though the data has not already been transferred to disk. *Id.*, col. 11:62-12:7. This liberates the processor from timing dependencies of prior disk subsystems and allows the processor "to dedicate its processing resources to increasing performance through more intelligent scheduling and data transfer control." *Id.*, col. 6:63-7:1. The use of cache memory also allows operations to be grouped together in cache and transferred in bulk to disk for

more efficient transfers. *Id.*, col. 6:55-60.

B. Representative Claim.

Claim 6 of the '857 patent is reproduced below with the disputed terms identified in bold.

> 6. In a disk memory system that comprises two data storage subsystems, each storing data records for said data processor, each of said data storage subsystems including a plurality of disk drives, a subset of said plurality of disk drives being configured into at least two redundancy groups, each said redundancy group consisting of at least two disk drives, wherein said two data storage subsystems are interconnected by a data link for exchanging control and data signals therebetween, a method of storing data records for at least one data processor in a fault tolerant mode, comprising the steps in one of said data storage subsystems of:
>
> > selecting, ***in response to the receipt of a stream of data records from said data processor***, available memory space in one of said redundancy groups to store said received stream of data records thereon;
> >
> > writing said received stream of data records and redundancy data associated with said received stream of data records in said selected available memory space in said one redundancy group;
> >
> > maintaining data indicative of the physical location of data records stored in said plurality of redundancy groups;
> >
> > transmitting, ***in response to writing a data record to said one redundancy group***, said written data record to the other of said data storage subsystems via said data link to maintain duplicate data records in both said data storage subsystems absent transmission of instructions to said data storage subsystem to create said duplicate data records by said data processor.

'857 patent, col. 13:57-14:20 (emphasis added).

C. Disputed Claim Terms To Be Construed.

1. "in response to writing a data record to said one redundancy group" (claim 6) / "responsive to writing a data record to one of said redundancy groups" (claim 11)

Sun's proposed construction	NetApp's proposed construction
in response to writing the data record to memory associated with the one redundancy group (claim 6) / responsive to writing the data record to memory associated with one of the redundancy groups (claim 11)	after and in reaction to writing a data record to a single redundancy group

Sun's proposed construction accurately describes how a write operation occurs in the '857

patent as clearly set forth in the specification. In contrast, NetApp's proposed construction contradicts the teaching of the specification – by excluding the only embodiment disclosed in the specification – and adds additional unsupported limitations to the claim.

a. Sun's Construction Conforms To The Specification, While NetApp's Construction Excludes The Only Embodiment Disclosed By The Specification.

In *Phillips*, the Federal Circuit stressed the importance of intrinsic evidence in construing claims, and in particular, the intrinsic evidence contained in the specification. Indeed, the Federal Circuit singled out the specification as "the single best guide to the meaning of a disputed term." *Phillips*, 415 F.3d at 1315 (*citing Vitronics*, 90 F.3d at 1582). Because the specification describes only one way of transmitting a data record between data storage subsystems, and because Sun's construction is consistent with that description, while NetApp's proposed construction contradicts that description, Sun's construction should be adopted.

The specification describes one embodiment for writing a data record to a redundancy group, including a plurality of disk drives. The write operation occurs in stages. First, the host processor transfers the data record to the first data storage subsystem. '857 patent, col. 11:51-54. Next, the first data storage subsystem writes the data record into its cache memory, which is associated with the redundancy group. *Id.*, col. 11:54-62. In response to writing the data record to cache memory, the first data storage subsystem transfers the data record to the second data storage subsystem. *Id.*, col. 11:62-12:1. Finally, the data record is written from cache memory to disk. *Id.*, col. 12:3-9. As emphasized in the '857 patent, the front-end operations, such as the write commands from the host processor to the data storage subsystems, are completely decoupled from the back-end operations, such as aggregating groups of data records in cache and transferring them to disk in bulk. *Id.*, col. 6:54-7:1. Therefore, as far as the host processor is concerned, a write operation is complete once the storage control units store the write in cache. This frees a host processor from waiting until data has been written to disk and allows it to perform other tasks to increase performance and improve efficiency. *Id.*, col. 6:63-7:1.

Sun's construction reflects these express and unambiguous teachings of the '857 patent. Under Sun's construction, representative claim 6 reads "transmitting, [in response to writing the

data record to memory associated with the one redundancy group], said written data record to the other of said data storage subsystems via said data link to maintain duplicate data records in both said data storage subsystems …" This comports exactly with the sequence of write operation steps described in the specification, namely, that a data record is transmitted from one data storage system to the other data storage subsystem ***after*** it has been written ***to cache***, but ***not*** necessarily ***after*** it is written ***to disk***. Thus, Sun's construction is consistent with both the language of claim 6 and the specification. Sun's construction also conforms to the patent's teaching that front end-operations are completely independent of the back-end operations, and that a write operation is considered complete from the perspective of a host processor once it has been written to cache memory.

In contrast, NetApp's construction contradicts the specification and excludes the preferred ***and only embodiment*** disclosed in the specification. NetApp's construction is on its face unclear as to whether the data record has to be written to disk.[1] However, NetApp's arguments clarify that it is in fact trying to limit the claims to a data storage subsystem that must wait until data is written ***onto disk*** before transferring the data to the duplicate data storage subsystem. *See, e.g.*, NetApp Op. Br., p. 15, n. 4. Not surprisingly, there is no support in the specification for NetApp's tortured interpretation of the claimed invention. Indeed, in direct contrast to NetApp's proposed construction, the specification clearly states that the data storage subsystem transfers a data record to a duplicate data storage subsystem ***before*** storing the data record to disk. *Id.*, col. 11:62-12:1. By excluding the preferred and only embodiment, NetApp's construction is inherently defective and must be rejected.

Although NetApp concedes that its construction excludes the preferred embodiment of the invention, it attempts to salvage the construction citing several cases. The case law NetApp cites

---

[1] NetApp's proposed construction says "after and in reaction to writing a data record to a single redundancy group." As explained in the specification, write operations to a redundancy group occur in stages and an operation is considered complete with respect to the host processor once it has been written ***to cache memory***. Therefore, "after and in response to writing data record to a redundancy group" could mean after and in response to writing the data record to cache memory.

is inapplicable and does not support a construction that would exclude the preferred ***and only*** embodiment of an invention. In each of the cited cases, the specification disclosed multiple embodiments, so a construction that excluded certain embodiments but not others was acceptable. *See Helmsderfer v. Bobrick Washroom Equipment, Inc.*, 527 F.3d 1379, 1382 (Fed. Cir. 2008); *Sinorgchem Co., Shandong v. Int'l Trade Comm'n*, 511 F.3d 1132, 1138 (Fed. Cir. 2007). However, in cases where there is only one embodiment disclosed, such as in the '857 patent, the case law soundly rejects constructions that would exclude the preferred and only embodiment. *Johns Hopkins Univ. v. CellPro, Inc.*, 152 F.3d 1342, 1355 (Fed. Cir. 1998) (holding "[a] claim construction that does not encompass a disclosed embodiment is thus 'rarely, if ever, correct and would require highly persuasive evidentiary support.') (citations omitted).

NetApp also argues that the '857 patent actually discloses an additional embodiment in which the order of steps does not matter. NetApp is wrong. NetApp points to the following language as evidence of this alleged second embodiment:

> Host processors 101, 121 transmit data record write and read requests to storage control units 105 and 107 in a conventional manner. The one of storage control units 105 and 107 that receives these requests (for example 105) communicates with its associated disk controller 112 and disk drives 109 to execute the write and read requests. In addition, ***in response to a received data record write command, storage control unit 105 transmits the received data record over data link 106 to storage control unit 107 to maintain identical virtual device images in both storage control units 105 and 107 and identical data records in data storage subsystems 1, 2.*** The system of FIG. 1 provides increased data availability over that which would be available if data storage subsystem 2 with disk drives 111, disk controller 113, storage control unit 107 and data link 106 were not provided.

'857 patent, col. 4:39-50 (emphasis added).

This brief description generally discusses and is wholly consistent with the single disclosed embodiment of a data record write operation, which is described in more detail in column 11 under the sub-heading "Data Record Write Operation." As explained in the excerpt above, the transmission of the data record from the first storage control unit to the second storage control unit is in response to receiving the data record from the host processor—not in response to writing the data record to disk. This is especially apparent from the sentence preceding the description of the transfer, which explains that the storage control unit first "receives" the data

record and then "communicates" it to disk. The "transfer" is described in the following sentence as occurring in response to receiving the data record command and not in response to communicating the data record to disk.

NetApp's construction also contravenes important features that the '857 patent teaches, namely, the importance of the back-end storage operations (*e.g.*, writes to disk) being completely independent from the front-end operations and the use of the cache memory to coalesce data record writes until a bulk transfer can be made. '857 patent, col. 3:59-4:10, 6:54-7:1. Under NetApp's construction, before transmitting a data record to a duplicate storage subsystem, a receiving data storage subsystem would have to wait until a large number of write operations were grouped together for a bulk transfer to disk. This would defeat the purpose of the cache memory and would subject the system to long periods without fault tolerance. For example, if the receiving subsystem failed before a bulk transfer to disk, the data records stored in its cache memory would be inaccessible to the host processor. Claims 6 and 11 both expressly state that the purpose of the claimed inventions is to store data records "in a fault tolerant mode." Because NetApp's construction is inconsistent with the context of the invention, it should be rejected. *Curtiss-Wright Flow Control Corp. v. Velan, Inc.*, 438 F.3d 1374, 1381 (Fed. Cir. 2006) (overturning District Court construction in part because "[a]ny construction to the contrary is not consistent with the overall context of this invention … as described in the specification.").

b. The Claim Language Supports Sun's Construction.

The claim language of the '857 patent further supports Sun's construction. The language of claims 6 and 11 say that the data record is transmitted "in response" to writing the data record to the redundancy group. As established above, the patent unambiguously explains that a write operation to a redundancy group occurs in several stages and ***always*** involves writing the data record ***to cache memory*** associated with the data storage subsystems ***before*** writing the record ***to disk***, and that a data storage subsystem transmits a data record to a duplicate subsystem ***after*** it has written it to ***cache***, but ***before*** it writes it ***to disk***. When Sun's proposed constructions are inserted into the body of the claims, the claims accurately reflect these clear teachings of the patent. *See, e.g.*, claim 6 ("transmitting, [in response to writing the data record to memory

associated with the one redundancy group], said written data record to the other of said data storage subsystems via said data link to maintain duplicate data records in both said data storage subsystems …").

In contrast, NetApp's proposed construction simply adds additional narrowing language (*i.e.*, "after and…") that is not found either in the specification or in the claims. Specifically, NetApp's construction replaces "in response to" with "*after and* in reaction to." The terms "in response to" and "in reaction to" are synonyms. Therefore, the net effect of NetApp's construction is to add the unsupported and narrowing language "after and" into the claims. The claim language, however, does not require the write operation be fully completed (*e.g.*, the claim does not say *after completing a write operation*), and the specification plainly describes the transfer happening ***after*** the data record has been written ***to cache*** memory associated with a redundancy group, but ***not after*** the data record has been written ***to disk***.

NetApp makes much of the fact that claim 6 says the data record transferred to the second data storage subsystem is a "written data record," arguing that a data record must be written to disk before it is transferred to the other data storage subsystem. In so saying, NetApp fails to recognize the multi-stage writing process described in the '857 patent and reflected in Sun's construction. As evident in Sun's construction, the data storage subsystem is in fact transmitting a "written data record" – *i.e.*, a data record that has been written to cache memory – to the other data storage subsystem.

Finally, NetApp argues that because there are so many claims (54) in the '857 patent, including claims using the language "in response to the *receipt of* said data record," that its construction is proper. First, NetApp cites no legal authority to support its absurd position that the number of claims justifies overriding the fundamental principles of claim construction rejecting constructions that do not encompass the single disclosed embodiment and recognizing the specification as "the single best guide to the meaning of a disputed term." *Phillips*, 415 F.3d at 1315 (*citing Vitronics*, 90 F.3d at 1582); *Johns Hopkins Univ.*, 152 F.3d at 1355.

Moreover, the claims cited by NetApp do not exclude the preferred embodiment. For example, claim 31 includes the term "means in said first data storage subsystem operable *in*

response to receipt of said data record over one of said channel interface paths for transmitting said data record over a data link transmission path to said second data storage subsystem." '857 patent, col. 37-41 (emphasis added). Clearly, this is consistent with the preferred and only embodiment because it does not exclude the possibility that the record is stored in cache memory. Indeed, an analysis of claim 32, which depends on claim 31, shows that the data record is, in fact, stored in cache "upon receipt," and further explains that a write complete signal is communicated to the host processor after the data record is stored in the cache memories of both data storage controls. *Id.*, col. 20:53-63. Thus, these claims are fully consistent with the preferred embodiment and provide no support for NetApp's proposed construction.

2. "in response to receipt of a stream of data records from said data processor" (claim 6) / "responsive to receipt of a stream of data records from said data processor" (claim 11)

Sun's proposed construction	NetApp's proposed construction
Sun contends that this phrase does not require construction because the phrase is clear on its face.	After and in reaction to the receipt of data records from a host processor.

There is no need to construe these phrases. The claim language itself is clear, and the patent specification does not further clarify the phrases.

As with the previous claim term, NetApp's construction simply replaces "in response to" with "*after and* in reaction to." And as with the previous claim term, there is nothing in the specification or claim language that justifies inserting an additional "after and" limitation into the claims. To the contrary, because the back-end storage operations (*e.g.*, writes to disk) are "completely independent" from the front-end operations in the '857 patent, the timing of back end operations (*e.g.*, selecting available memory space) is decoupled from the timing of front end operations (*e.g.*, receiving data records). *Id.*, col. 3:59-4:10, 6:54-7:1. In other words, the back-end operations need not wait until front-end operations are fully completed in order to begin. *Id.*, col. 6:63-71. NetApp's construction, which contemplates coupling the timing of back-end and front-end operations, contradicts this teaching of the specification. Accordingly, NetApp's construction should be rejected.

IV. U.S. PATENT NO. 5,749,095

A. Technology Background.

The invention of the '095 patent allows computer systems with multiple processors to accelerate performance. Declaration of Dr. Donald Alpert In Support Of Sun's Responsive Claim Construction Brief ("Alpert Decl."), ¶ 10; '095 patent, col. 3:5-8. Specifically, the '095 patent teaches a breakthrough technique for rapidly writing data from a processor to system memory in a multiprocessor system without suffering the delays that degraded computer system performance in the prior art. '095 patent, Abstract.

The application of the fast write protocol is described within the '095 patent in the context of a multiple processor computer with a distributed shared memory architecture. By distributing the main memory among each of the processors or nodes in a multiprocessor system, several typical bottlenecks can be avoided. '095 patent, col. 4:19-22. An example of a computer system with four processors implementing a distributed shared memory is illustrated in Figure 1A from the '095 patent:

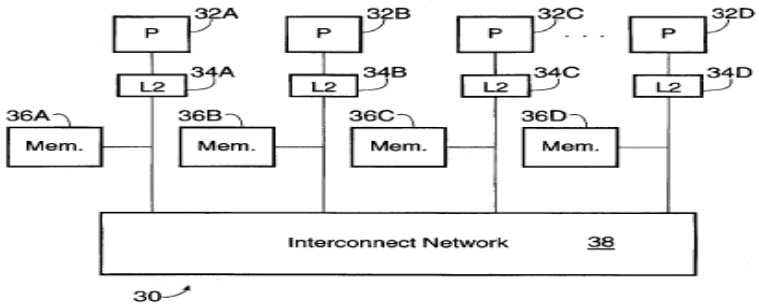

Fig. 1A

'095 patent, Fig. 1A.

Within this example, the main memory is comprised of four devices (36A, 36B, 36C, and 36D). Each is associated with a processor (32A, 32B, 32C, and 32D). '095 patent, col. 9:66 – col. 10:4. Each processor also has an identified cache memory (34A, 34B, 34C, and 34D) and the capability of communicating with other processing nodes. '095 patent, col. 10:1-4, col. 31:19-20. In the parlance of the '095 patent, a "node" is formed, at a minimum, by a processor and its associated memory. *Id.*, col. 18:1-6, col. 31:17-18.

The multiple memories 36 form a distributed shared memory where each location of the

memory is referenced by an address. *Id.*, col. 10:4-6. Using an address, any node can reference any location within any one of the memory devices owned by another node. *Id.*, col. 9:31-35. Thus, it is common that a processing node will read data from and write data to a memory location within another processing node. *Id.*, col. 23:4-65.

The drawback of the distributed shared memory is encountered when a processing node wants to ***modify data*** stored in a memory location associated with ***another*** processing node. *Id.*, col. 4:41-46. To perform a "coherent" operation related to modifying data in another node, the processing nodes communicate to ensure the request is reflected in the cache of each processing node holding a copy of the data. *Id.*, col. 3:41-44, col. 18:16 - col. 19:12. There is a delay associated with these "coherency operation" communications that can result in degradation of the system performance. Alpert Decl., ¶ 28.

This issue is amplified when each processing node within the system seeks to initiate more than one coherency operation at the same time. '095 patent, col. 20:31-34. Each time a coherency operation is requested, resources within the processor of the requesting node are normally occupied until the operation is complete. *Id.*, col. 4:46-55. The delay involved in performing the coherency operations eventually may tie up all processor resources. Alpert Decl., ¶ 31. This occupation of processor resources results in poorer performance. *Id.*

The '095 patent teaches a fast write protocol to avoid severe degradation of processor performance. '095 patent, col. 4:62-64. A fast write is performed by sending the data from the processor prior to or concurrently with the performance of the related coherency operation. *Id.*, col. 5:5-8, col. 27:25-28, col. 31:36-41; Alpert Decl., ¶ 33. By performing the coherency operation after writing the modified data and making it accessible, performance is improved by freeing processor resources more rapidly. '095 patent, col. 7:58-50, col. 28:29-33, col. 31:3-7, col. 31:36-41, col. 31:41-44.

Using specific encoding to request a fast write, updated data can be provided immediately by the processor. *Id.*, col. 28:46-50. This frees the resources occupied by the modified data within the processor and enables the new coherency operations. *Id.*, col. 28:28-33. However, in order to move on to new operations, the processor must be guaranteed that the coherency

operation will be completed by the system. *Id.*, col. 28:53-57, col. 31:38-41, Fig. 15 (item 324). Thus, the '095 patent teaches the coherency operation may be not omitted and the processing node must be guaranteed the coherency will be completed. *Id.*, col. 27:16-23, col. 27:41-49.

B. Representative Claim.

Claims 1, 11 and 17 are the three independent claims asserted by Sun against NetApp. Representative claim 1 is reproduced below with the disputed claim term in bold.

> 1. A method for performing write operations in a multiprocessing computer system, comprising:
> initiating a write operation by a processor within a local processing node of said multiprocessing computer system;
> performing a coherency operation to at least one remote processing node in response to said write operation;
> **completing said write operation within said local processing node** prior to completion of said coherency operation if said write operation includes a specific predefined encoding; and
> **completing said write operation within said local processing node** subsequent to completion of said coherency operation if said write operation includes an encoding different than said specific predefined encoding.

'095 patent, col. 31:63 – col. 32:10 (emphasis added).

C. Disputed Claim Term To Be Construed.

1. "Completing [a] write operation within [a] local processing node" (claim 1) / "complete [a] write operation with respect to [a] processor" (claims 11, 17)

Sun's proposed construction	NetApp's proposed construction
Claim 1: Sun contends this phrase does not require construction because the phrase is clear on its face. However, if the Court decides to construe this phrase, the phrase should be construed to mean: "data for the write operation (1) is provided to subsequent read operations within the local processing node to the same address as the write operation and (2) is or will be coherent within distributed shared memory."	Claim 1: "transferring the write data from an initiating processor to a system interface."
Claims 11 & 17: Sun contends this phrase does not require construction because the phrase is clear on its face. However, if the Court decides to construe this phrase, the phrase should be construed to mean: "data for the write operation (1) is provided to	Claim 11 & 17: "transferring the write data from an initiating processor to a system interface."

subsequent read operations by a processor to the same address as the write operation and (2) is or will be coherent within distributed shared memory."	

This claim language is clear on its face and does not need to be construed because, as established below, one of ordinary skill in the art – including NetApp's own expert – knows precisely what the language means. Moreover, NetApp's construction conflicts with the language of claim 1, and improperly reads the "system interface" into all of the claims even though it is not part of claims 1 or 17.

If the Court decides to construe these terms, the Court should adopt Sun's constructions because they are based on the intrinsic record and comport with the well-accepted meaning of these terms to those of ordinary skill in the art. The claim language and specification inform a person of ordinary skill in the art that the write operation is complete once the processor initiating the write operation can successfully read from the address of the write operation and is guaranteed that coherency will be performed. Alpert Decl., ¶ 39.

a. The Claim Language Is Clear On Its Face To Person of Ordinary Skill In The Art.

The meaning of completing a write operation is clear to a person of ordinary skill in the art. Alpert Decl., ¶ 39. A person of ordinary skill in the art recognizes the term "complete" in the context of a write operation requires the written data be capable of being read. *Id.*, ¶¶ 39-40. The determination of when a write operation is complete was settled prior to the '095 patent.

Well-regarded publications in the field state a write operation is complete in a system when "the value returned on a LOAD instruction is always the value given by the latest STORE instruction with the same address." Alpert Decl., ¶ 43. This means that a write operation is complete if any portion of the system will read (LOAD) the data of the write (STORE) operation from the same address as the write operation. *Id.*, ¶ 40. In the context of a node within a system, a write operation is complete if that node can read the data of the write operation from the same address as the write operation and the node is guaranteed that the write will be coherent within the system. *Id.*, ¶ 40. This is the plain meaning of the claim language to one of ordinary skill in the

art. *Id.*

This common understanding of the claim language has been recognized by NetApp's expert, Dr. DuBois, in his published works. For example, in his co-authored paper, *Correct memory operation of cache-based multiprocessors*, Dr. DuBois described the coherence of a write in a system as determined by whether "the value returned on a LOAD instruction is always the value given by the latest STORE instruction with the same address." Alpert Decl., Ex. D, at 235, Scheurich, C. and Dubois, M., *Correct memory operation of cache-based multiprocessor*, Proceedings of the 14th Annual international Symposium on Computer Architecture (Pittsburgh, Pennsylvania, June 2-5, 1987). Directly to the point, Dr. DuBois defined "performing with respect to a processor" as:

> A STORE by processor i is considered performed with respect to processor k at a point in time when a subsequently issued LOAD to the same address by processor k returns the value by a STORE …

Id. at 236. This is the same understanding held by Sun's expert, Dr. Alpert. Alpert Decl., ¶¶ 44-46.

Thus, Dr. Dubois' work outside this litigation tracks the plain meaning of this claim language – which is mirrored in the construction offered by Sun. Dr. DuBois' prior statements completely belie his current litigation-driven "opinion" that these claim terms "do not have well understood meaning in the art." Dubois Decl., p. 7. Therefore, the overwhelming evidence – supported by the work of both parties' experts – demonstrates the existence of an accepted plain and ordinary meaning for this claim term such that no construction of this claim term is required.

b. NetApp's Construction Is Inconsistent With Claim 1.

The language "completing said write operation within said local processing node" is recited twice within claim 1. '095 patent, col. 32:3-4, col. 32:7-8. The first recitation allows the write operation to be completed "prior to" completion of the coherency operation. *Id.*, col. 32:3-6. The second recitation differs by requiring the write operation to be completed "subsequent to" completion of coherency operation. *Id.*, col. 32:7-10.

The specification is replete with examples of write operations that are completed "subsequent to" the completion of the coherency activity. These examples include write

operations where updated ***data is actually transferred*** to the other caches in the system and alternatives where ***no data is transferred*** at all. *Id.*, col. 3:44-50. A write operation may be completed without transferring data to other caches in the system by ensuring that any outdated copies are "invalidated." *Id.*, col. 17:57-67. This type of write operation is detailed in Figures 4 and 5. *Id.*, Fig. 4, col. 19:4-12, Fig. 5, col. 20:21-23. Since the write data associated with these operations is not to be stored in any remote cache, ***the processor need not transfer the data to the system interface***, as would be required by NetApp's construction. Alpert Decl., ¶¶ 49-50. Yet, it is clear that these operations complete the write operation with respect to the processor subsequent to a coherency operation and fall within the last limitation of claim 1. Alpert Decl., ¶¶ 49-50. Because NetApp can find no justification for excluding from the scope of the claims these write operation embodiments that are expressly disclosed in the specification, NetApp's construction of this claim term cannot be squared with the second recitation of this claim term in claim 1. Alpert Decl., ¶¶ 49-50. This flaw in NetApp's construction requires that NetApp's construction be rejected.

c. NetApp's Construction Impermissibly Reads The System Interface Into Claims 1 and 17.

NetApp's construction also is incorrect because it incorporates the "system interface" described in the specification, and expressly recited in claim 11, into independent claims 1 and 17. However, claims 1 and 17 do not recite a system interface requirement and thus are of different scope.

NetApp's brief provides an excerpt of claim 11 and emphasizes the recitation of a system interface. NetApp Op. Br., at 19. NetApp relies on claim 11 as "representative" of the other claims being considered. *Id.* Oddly, though, NetApp omits any excerpt from claims 1 and 17 and fails to mention to the Court that the words "system interface" do not appear in those claims.

Moreover, rather than supporting NetApp's position, the language of claim 11 establishes that a "system interface" limitation ***cannot*** be read into the other claims. Claim 11 reads in part:

> wherein said system interface is configured to complete said write operation with respect to said processor . . .

'095 patent, col. 32:41-43. Thus, claim 11 includes an express limitation providing that the "system interface" "completes" the write operation. *Id.* Neither this limitation, nor any recitation of a "system interface," is included in independent claims 1 or 17 (or their respective dependent claims).

NetApp's attempt to impose a "system interface" limitation on claims 1 and 17 violates many rules of claim construction. First, as acknowledged by NetApp's brief, a claim without a "system interface" is broader than a claim with such a limitation. Sun agrees, and the Federal Circuit holds that where some claims are broad and others are narrow, the narrow claim limitation cannot be read into the broader claims. *Uniroyal*, 837 F.2d at 1054-55.

Second, NetApp's proposed construction of this claim term would render superfluous the above-quoted limitation of claim 11 as it would result in the same limitation being imposed twice. This is prohibited. *Innova/Pure*, 381 F.3d at 1119; *Merck*, 395 F.3d at 1372; *Phillips*, 415 F.3d at 1324.

Third, when different claims use different language – such as the presence of the above-quoted limitation in claim 11 and its absence in claims 1 and 17 – a presumption exists that the claims have different meanings. *CAE Screenplates, Inc. v. Heinrich Fielder GmbH & Co. KG*, 224 F.3d 1308, 1317 (Fed. Cir. 2000); *Comark Communications, Inc. v. Harris Corp.*, 156 F.3d 1182, 1187 (Fed. Cir. 1998). Here, NetApp attempts to conflate claims 1, 11 and 17 on this point in violation of this rule.

Fourth, it is a fundamental rule of claim construction that the Court is not authorized to read into a claim a limitation that is lacking, as is the case with the "system interface" in claims 1 and 17. *Renishaw PLC v. Marposs Societá Per Azioni*, 158 F.3d 1243, 1249 (Fed. Cir. 1998); *Zenith Lab. Inc. v. Bristol-Myers Squibb Co.*, 19 F.3d 1418, 1422 (Fed. Cir. 1994). Under each of these four rules of claim construction, NetApp's construction is improper and must be rejected.

While NetApp's brief is short on discussion of the actual claim language – and its legal impact on claim construction – NetApp's brief is long on its discussion of the specification. Throughout, NetApp argues that the system interface has the "central role" within the described embodiments. NetApp Op. Br., at 20. In fact, NetApp argues the "heart of the alleged invention

of the '095 patent is a 'system interface'…." NetApp Br. at 18. This argument is contrary to established law. *Allen Eng'g Corp. v. Bartell Indus.*, 299 F.3d 1336, 1345 (Fed. Cir. 2002) (repeating the prohibition against substituting an alleged ***heart of the invention*** for the claims); *Phillips*, 415 F.3d at 1314 (recognizing the importance of the language of the claim).

NetApp's argument also is contrary to the specification and to the understanding of those of ordinary skill in the art. From the perspective of those of ordinary skill in the art, the '095 patent describes the fast write protocol broadly and in a manner that is not limited to a hardware embodiment. Alpert Decl., ¶¶ 51-53. Indeed, the Summary of the Invention section of the specification is clear that the fast write protocol falls within the scope of the invention without regard to the "system interface" hardware in claim 11. '095 patent, col. 5:29-51. In fact, a person of ordinary skill in the art would recognize that the alternative inventions claimed in each of the three independent claims track each of the last three paragraphs of the Summary of the Invention. Alpert Decl., ¶ 52; '095 patent, col. 5:39 – col. 6:12.

The specification describes the manner in which the inventive fast write is executed. '095 patent, col. 5:29-32. The specification details the invention as a method and in broad terms. *Id.*, col. 5:39-51. After discussing the invention in terms of the method, the specification describes a hardware configuration performing a fast write, and does so without any mention of the system interface. *Id.*, col. 5:66 – col. 6:12. Thus, the implementation of the fast write method is not tied to any one hardware platform or to implementing a system interface. Alpert Decl., ¶ 53.

NetApp's argument is further weakened by the specification's repeated confirmation that the invention may be implemented in alternate forms. Throughout the specification, the '095 patent confirms that the invention covers "all modifications, equivalents and alternatives falling within the spirit and scope of the present invention," and that the invention is not "limit[ed] … to the particular form disclosed." '095 patent, col. 7:9-18; col. 31:54-59. Indeed, the specification states in the context of Figure 2 and the system interface that "[i]t is noted that alternative specific implementations of each SMP node 12 of FIG. 1 are also possible." '095 patent, col. 11:20-21.

d. If A Construction Is Necessary, The Specification Supports Sun's Proposed Construction.

The specification uniformly recognizes that write operations require that coherency operations be performed either prior to or after performing the write. '095 patent, col. 19:4-12, col. 27:17-23. The specification simply does not teach performing a write to the memory of a remote node without ensuring at some point that the write is coherent among each cache in the system. Thus, the specification supports Sun's proposed construction that completing the write operation with respect to a processor requires that data for the write operation is or will be coherent within distributed shared memory.

V. U.S. PATENT NO. 6,873,630

A. Technology Background.

The '630 patent is directed towards an Ethernet network architecture that enables Ethernet networks to operate at increased transmission rates of multiple gigabits per second. '630 patent, col. 1:6-9, 61-65; Declaration of Greg Ennis In Support of Sun's Responsive Claim Construction Brief ("Ennis Decl."), ¶ 15. The invention of the '630 patent proved to be such an improvement for Ethernet networks that it was incorporated into the IEEE 802.3 Ethernet network standard. *Id.*

The invention of the '630 patent enables increased transmission rates by dividing communications into smaller parts and distributing the smaller parts among a plurality of logical channels. '630 patent, col. 2:28-47; Ennis Decl., ¶ 19. By sending the communication across multiple logical channels rather than a single channel, the overall data transmission rate is increased to approximately the sum of the data transmission rates of each of the multiple logical channels, rather than being limited to the data transmission rate of a single channel. *Id.*

Some embodiments of the invention involve first receiving from a computer's MAC module a data stream communication. '630 patent, col. 3:13-15, col. 20:40-42; Ennis Decl., ¶¶ 19-20. The data stream is comprised of Ethernet frames and additional control information

associated with the Ethernet frames that is necessary for properly transmitting the frames.[2] Ennis Decl., ¶ 19. This data stream is then divided into multiple portions, each of which is transmitted across one of many logical channels. *Id;* '630 patent, col. 2:28-47, col. 9:24-33. When these portions of the data stream reach the destination computer, they are combined into a single data stream and passed to the destination computer's MAC layer. Ennis Decl., ¶¶ 20-21; '630 patent, col. 2:33-35.

B. Representative Claims.

The '630 patent includes nineteen independent claims—claims 1, 3, 15, 22, 25, 27, 30, 34, 42, 45, 48, 52, 59, 62, 63, 72, 73, 89 and 123—of which claims 3, 15, 22, 25, 30, 45, 48, 52, 62, 73 and 89 are being asserted against NetApp. Representative claim 3 is reproduced below with instances of the two disputed claim terms identified in bold.

> 3. A method of transmitting a communication from a first network entity to a second network entity, wherein the first network entity and the second network entity are coupled to a communication medium, comprising:
>
> receiving a communication at a distribution module of a network interface device from a medium access control module across a first interface, wherein said distribution module is configured to distribute **portions of said communication** among a plurality of communication channels;
>
> distributing **elements of said communication** into multiple portions;
>
> sending a first **portion of said communication** on a first channel established on a first communication medium coupled to said first network entity and said second network entity; and
>
> sending a second **portion of said communication** on a second channel established on a second communication medium coupled to said first network entity and said second network entity

'630 patent, col. 15:8-28.

[2] In its brief, NetApp incorrectly asserts that the "highest level" of a communication is a frame. NetApp Op. Br., at 28, ll. 17-20. Instead, the highest level of a communication is its data stream, comprising frames and associated data. Ennis Decl., ¶ 19.

C. Disputed Claim Terms To Be Construed.

1. "portion [of a] communication"

Sun's proposed construction	NetApp's proposed construction
No construction necessary, the term is clear on its face.	The fraction or portion of a frame carried by one channel.

The phrase "portion [of a] communication" contains no technical terms or terms of art, and consists entirely of plain English words. Ennis Decl., ¶ 25. Accordingly, the phrase should be given its plain and ordinary meaning. *Ecolab, Inc. v. Envirochem Inc*., 264 F.3d 1358, 1366 (Fed. Cir. 2001) ("[W]e presume that the terms mean what they say. In other words, we follow the general rule that terms in the claim are to given their ordinary and accustomed meaning."); *Johnson v. Worldwide Assocs., Inc. v. Zebco Corp*., 175 F.3d 985, 989 (Fed. Cir. 1999) (there is a "heavy presumption in favor of the ordinary meaning of claim language"); *Renishaw*, 158 F.3d at 1249 ("when a claim is expressed in general descriptive words, we will not ordinarily limit the terms").

Moreover, should the Court conclude a construction is appropriate, the Court should reject NetApp's construction because it is incorrect for the reasons identified below.

a. NetApp Improperly Equates A "Communication" With One Frame.

NetApp's proposed construction of a "portion [of a] communication" as "the fraction or portion of a frame carried by one channel" is built on a major, but entirely erroneous, premise—that the ***entire*** "communication" is a ***single*** frame. The specification, however, leaves no doubt that a communication is simply a ***stream of data***, which may comprise ***more than one frame***. Ennis Decl., ¶ 26. The specification states: "In one embodiment of the invention the communication is divided for transmission across multiple channels . . . in this embodiment the individual bytes of ***each frame***, or packet, ***of the communication*** are separated and sent across one of the channels in a round-robin fashion." '630 patent, col. 2:36-41 (emphasis added). Thus, a "communication" clearly can comprise more than one frame. Ennis Decl., ¶¶ 27-28.

The fact that a "communication" consists of a stream of data, rather than simply being a single frame, is confirmed throughout the specification. For example, the specification teaches that "by distributing or striping a data stream across multiple channels, the data stream can be transmitted at substantially the sum of the individual channels." '630 patent, col. 4:26-30. Thus, the patent teaches that it is a ***data stream*** that is "distributed" across multiple channels. The specification teaches that in a preferred embodiment this "distribution" is carried out by a distributor/collector that "performs a distribution function for ***a communication*** sent from an attached computer system in order to disseminate ***portions of the communication*** across the multiple logical channels. When receiving ***a communication***, however, distributor/collector 100 collects data from the multiple channels to ***re-assemble*** a single ***data stream*** to pass to the attached network entity . . ."[3] *Id.*, col. 4:57-65 (emphasis added).

In other words, the "distributor" takes an outgoing ***communication*** from an attached computer system and disseminates its portions across the multiple logical channels, while the "collector" takes an incoming distributed communication and "reassembles" it into a single ***data stream***. Ennis Decl., ¶ 28. Thus, as is readily appreciated by one of ordinary skill in the art, the specification uses the terms "communication" and "data stream" interchangeably. *Id.*

Moreover, the communication (*i.e.*, data stream) comprises not only multiple frames (or packets), but also additional elements necessary for transmission. *Id.*, ¶¶ 29-34. This is shown in Figure 4 of the '630 patent, which depicts the transfer of a 64 byte packet followed by multiple 65 byte packets, with each frame accompanied by idle codes (IPG) and preamble codes (PA). '630 patent, col. 12:42-13:13; Ennis Decl., ¶¶ 29-30.

All of these disclosures leave no doubt as to the plain meaning of "portion [of a] communication." One of ordinary skill in the art, upon reviewing the specification and the claims, would have no trouble ascertaining that "portions [of a] communication" may and do encompass portions of a multi-frame data stream, and not just portions of a single frame within

[3] As noted by NetApp, this key passage is the only place in the specification where the term "portion of a communication" appears.

that data stream, as proposed by NetApp. *Id.*, ¶¶ 26-28.

This understanding is fully consistent with the parties' agreed upon construction of "distribution module," which NetApp misinterprets as supporting its narrow construction of "portion [of a] communication." The agreed upon construction of "distribution module" states a distributor is "a module that divides across multiple logical channels an Ethernet frame received from a MAC module." Each such frame arrives at the "distributor" as part of the data stream comprising the communication shown in Fig. 4. Ennis Decl., ¶¶ 29-31. The distributor operates on the communication frame by frame, but the resulting "portion of the communication" does not consist of only a part of a frame (as NetApp's construction incorrectly asserts), but rather consists of a subset of the original data stream, including both multiple portions of multiple frames and necessary inter-frame data, such as preamble and idle codes. '630 patent, Figs. 5A-D, col. 13:15-53; Ennis Decl., ¶¶ 31-33.

b. NetApp's Construction Eviscerates Distinctions Between Claims.

Equating a "communication" with a "frame" also is incorrect because some of the claims of the '630 patent expressly recite "frames." For example, claim 62 recites, in part, "a medium access control module configured to communicate a first ***frame***…." '630 patent, col. 22:56-57 (emphasis added). "In the absence of any evidence to the contrary, we must presume that the use of … different terms in the claims connotes different meanings." *CAE Screenplates*, 224 F.3d at 1317; *Comark Communication*, 156 F.3d at 1187 ("There is presumed to be a difference in meaning and scope when different words or phrases are used in separate claims."). Because the claims use "frame" in some places and "communication" in other places, those two terms are presumed to have different meanings. However, NetApp's proposed construction eviscerates any distinction in meaning between these two claim terms.

NetApp's proposed construction also ignores that dependent claims 2, 16, 44, 39 and 53 limit the term "communication" to an "Ethernet frame" or "packet." '630 patent, col. 15:5-6, 16:54-55, 20:31-34, 21:56-60. For example, claim 44 depends on claim 3 and recites "wherein said communication is an Ethernet frame." *Id.*, col. 20:31-32. As a matter of law, because the

narrower dependent claims limit "communication" to a "frame," the term "communication" in the broader independent claims cannot be so limited. *Phillips*, 415 F.3d at 1315 ("the presence of a dependent claim that adds a particular limitation gives rise to a presumption that the limitation in question is not present in the independent claim"); *Liebel-Flarsheim Co. v. Medrad, Inc*., 358 F.3d 898, 910 (Fed. Cir. 2004) ("As this court has frequently stated, the presence of a dependent claim that adds a particular limitation raises a presumption that the limitation in question is not found in the independent claim.").

c. NetApp Mischaracterizes The Intrinsic Record.

NetApp suggests its construction is bolstered by the part of the specification stating the term frame "generally refer[s] to the unit of information received from or sent to a MAC layer from a physical layer device." '630 patent, col. 5:48-51. But nowhere does the specification equate this single unit of information with the term "communication." In fact, as established above, the specification is clear the term "communication" is not so limited. Ennis Decl., ¶¶ 26-28.

NetApp cites the specification's discussion of a "mini-frame" as allegedly supporting its proposed claim construction. NetApp Op. Br., p. 30. In fact, elsewhere in its brief, NetApp defines "portion of a communication" as a "mini-frame," thereby reading the term mini-frame into every claim that contains the term "portion of a communication." *Id*., p. 34. However, the term "mini-frame" is not present in any of the independent claims, appearing only in dependent claims 53-56. The Federal Circuit has made clear that reading limitations from narrower dependent claims into independent claims always is improper. *Amgen Inc. v. Hoechst Marion Roussel, Inc.*, 314 F.3d 1313, 1326 (Fed. Cir. 2003) ("Our court has made clear that when a patent claim does not contain a certain limitation and another claim does, that limitation cannot be read into the former claim in determining either validity or infringement."); *Phillips*, 415 F.3d at 1315 ("the presence of a dependent claim that adds a particular limitation gives rise to a presumption that the limitation in question is not present in the independent claim").

NetApp's reference in footnote 7 of its brief to certain instances of the specification's use of the word "portion" fails for the same reason. Specifically, each instance cited by NetApp

involves the use of a "mini-frame." '630 patent, col. 2:48-51, col. 5:51-53, col. 12:19-23.

NetApp also cites the specification at column 8, lines 24-43 as purportedly "explain[ing] the shortcomings of an approach where something other than 'fractions or portions of [a] frame' are used as the 'portions of the communication.'" NetApp Op. Br., p. 31. But neither phrase NetApp puts in quotation marks is in fact present in the cited passage. Instead, the cited passage first describes a prior art link aggregation technique, where the data stream was split above the MAC level. '630 patent, col. 8:23-31. The specification then describes an embodiment of the '630 invention as "striping [] network data at a lower level in the network protocol stack" than the prior art. *Id.*, col. 8:33-39. This striping of network data is done "with the contents of individual MAC frames or packets." *Id.* at col. 8:39-42.

NetApp does not explain how any of this disclosure supports its contention that a "portion of a communication must be a portion of a frame." In fact, the very passage cited by NetApp shows that NetApp's construction is untenable. The specification uses the term "network data" to describe data that is distributed to multiple channels—a communication that is distributed into portions. This network data certainly contains MAC frames, the contents of which are striped, but this does not mean that the network data is one single frame or that a portion of that data is a fraction of one single frame. To the contrary, and as established above, the specification, including the very passage cited by NetApp, teaches that a communication is a data stream (comprising network data arranged in MAC frames) and that a portion of that communication is a portion of that data stream. Ennis Decl., ¶¶ 26-36.

NetApp also contends that "claims of the '630 patent that refer to a 'portion of a communication' confirm [NetApp's construction]" because the claims refer to a "'communication' as being received or sent to a 'media access control module.'" NetApp Op. Br., p. 31. This is incorrect. While some (but not all) claims require the communication be received from a MAC module (*e.g.*, claim 3), there is nothing in the claim language requiring the communication be limited to a single frame. In fact, the claim language relied upon by NetApp is fully consistent with the plain meaning of "communication"—a data stream. Ennis Decl., ¶¶ 34-35. In claim 3, for example, the data stream is received at a distribution module from a MAC

module (col. 15:12-14) and is operated on by a distribution module that distributes its portions among a plurality of communication channels (col. 15:15-17). As reflected in the parties' agreed construction of "distribution module," the module operates on the frames within the data stream. As established above, the distributor operates on all the frames in the data stream and on other data within the data stream (such as idle codes), and outputs a "portion of the communication" consisting of multiple portions of multiple frames and necessary inter-frame data.

Amazingly, while NetApp initially concedes the patent specification contains no disclaimers or express limitations on the scope of the claims – and, in fact, makes clear that the enabling embodiment described in the specification is just illustrative – NetApp then asserts that Federal Circuit law regarding express limitations in the specification, and descriptions of what constitutes "the invention," should apply. NetApp Op. Br., pp. 29, 31. However, as NetApp itself initially noted, the factual predicate for the application of such precedent is lacking. Indeed, NetApp has not overcome the heavy presumption in favor of the ordinary meaning of this claim term. In order to vary the ordinary meaning of a term, a patentee must *unambiguously indicate* any special meaning of a claim term in the specification. *York Prods., Inc. v. Central Tractor Farm & Family Ctr.*, 99 F.3d 1568, 1572 (Fed. Cir. 1996) (a claim term takes on its ordinary meaning unless there is an "*express intent*" to impart a novel meaning to the term) (emphasis added); *Teleflex*, 299 F.3d at 1325 (deviation from a term's plain and ordinary meaning requires "a clear disavowal of claim scope"). No such express intent or clear disavowal of claim scope exists here.

d. NetApp Improperly Requires Every Portion Of A Communication To Be Carried By One Channel.

NetApp's inclusion of the added limitation "carried by one channel" in its proposed construction of "portion [of a] communication" also is incorrect. As an initial matter, nothing in the phrase "portion [of a] communication" refers to channels, the carrying of the communication across channels, or the number of channels over which a portion of a communication must be carried. "We know of no principle of law which would authorize us to read into a claim an element which is not present…." *Renishaw PLC*, 158 F.3d at 1249; *Zenith Lab.*, 19 F.3d at 1422

("It is improper to read a limitation into a claim wholly apart from any need to interpret what the patentee meant by particular words or phrases in the claim.").

Furthermore, the addition of this added limitation is unnecessary as the surrounding claim language already calls for at least some portions of a communication to be sent over specified channels. '630 patent, col. 15:20-25. As a matter of law, the proper construction of a claim term cannot be derived by pulling the surrounding claim language into the claim term, so as to create a tautology within the claim. *See AllVoice Computing PLC v. Nuance Communications, Inc.*, 504 F.3d 1236, 1247-48 (Fed. Cir. 2007) (refusing to define a limitation to encompass a second limitation in the same claim); *Mangosoft, Inc. v. Oracle Corp.*, 525 F.3d 1327, 1330-31 (Fed. Cir. 2008) (rejecting a proposed construction that would have rendered the disputed claim term "superfluous" and "ascribe[d] no meaning to the [disputed claim term] not already implicit in the rest of the claim").

2. "element [of a] communication"/"element [of a] communication portion"/"elements"/"element [of a] portion"

Sun's proposed construction	NetApp's proposed construction
No construction necessary, the term is clear on its face.	A portion (e.g., a byte) of a mini-frame that is individually encoded for transmission across one of a plurality of logical channels, where a mini-frame is a fraction or portion of a communication received from or sent to a media access control layer from a physical layer device and that is carried by one channel.

The word "element" is a simple English word. That simple word, and the claim phrases using that word, require no construction as they are well-understood by one of ordinary skill in the art. Ennis Decl., ¶¶ 23-24.

NetApp responds to this simple word with a litigation-driven, fifty-five word construction that amounts to a house of cards. In doing so, NetApp takes its erroneous construction of "portion of a communication" and embellishes it by impermissibly incorporating several limitations from various dependent claims and preferred embodiments.

Here is how NetApp stacks the cards. According to NetApp: "Because an 'element' is

part of a 'portion of a communication,' and a 'portion of a communication' is a 'mini-frame,' it follows that an 'element of a communication' must be part of a 'mini-frame.'" NetApp Op. Br., p. 34. Note that having previously argued in its brief that a "portion of a communication" is a "fraction or portion of a frame carried by one channel," NetApp now, in this construction, further asserts that a "portion of a communication" must be a "mini-frame." However, no "mini-frame" limitation is present in any independent claim, and instead only appears in dependent claims 53-56. Accordingly, as established above in connection with the term "a portion of a communication," neither that claim term nor this claim term can be limited to use of "mini-frames." *Amgen*, 314 F.3d at 1326; *Phillips*, 415 F.3d at 1315.

NetApp places another card on its already unstable house by requiring that the "portion of a mini-frame" be "individually encoded for transmission across one of a plurality of logical channels." The concept of encoding, however, also is claimed in dependent claims, and accordingly, cannot be used to limit the independent claims. *Id.*

As its final card, NetApp requires, without a word of explanation, that the mini-frame be "a fraction or portion of a communication received from or sent to a media access control layer from a physical layer device." There is no support for this limitation. Accordingly, like any house of cards, NetApp's proposed construction must fall.

a. NetApp Improperly Reads "Mini-Frames" Into Most Independent Claims.

Having incorrectly construed "portion of a communication" as a "fraction or portion of a frame carried by one channel," NetApp uses the same faulty logic to contend that "element of a communication" is a portion of a "mini-frame," which NetApp asserts "is substantively equivalent to how the term 'portion of a communication' should be understood." NetApp Op. Br., p. 34. Essentially, NetApp argues that "element of a frame" should be construed as a "portion of a portion of a frame carried by one channel," with some additional limitations. However, as demonstrated above, neither the claims nor the specification limit a "portion of a communication" to a portion of a frame carried by one channel. To the contrary, as established above, a portion of a communication may contain portions of multiple frames and other data

necessary for transmission. For this reason alone, NetApp's construction must be rejected.

Moreover, NetApp's construction injects the term "mini-frame" into every claim that refers to "elements." By NetApp's own count, these are claims 3, 5, 8, 13-15, 21-23, 45-48, 50, 52, 73, 76, 79, 89, 113, 114, 116 and 117. NetApp Op. Br., p. 33. Notably missing from this long list are dependent claims 53-56, the only claims that actually recite the term "mini-frame." This is another instance of NetApp attempting to read limitations from narrow dependent claims into almost every independent claim. As established above, this is improper. *Phillips*, 415 F.3d at 1315; *Amgen*, 314 F.3d at 1326; *see also Comark Communications*, 156 F.3d at 1187 ("There is presumed to be a difference in meaning and scope when different words or phrases are used in separate claims."); *CAE Screenplates*, 224 F.3d at 1317 (same).

Furthermore, even if NetApp were correct (which it is not) to limit "portion of a communication" to a "mini-frame," NetApp's construction of "mini-frame" is deficient. As NetApp correctly notes, the specification defines a "mini-frame" as a "portion of the frame carried by one channel." '630 patent, col. 2:48-51. NetApp's proposed construction, however, inexplicably differs from the specification by requiring a mini-frame to be "a fraction or portion of a communication *received from or sent to a media access control layer from a physical layer device* and that is carried by one channel." The italicized portion of NetApp's construction is not part of the definition of "mini-frame" in the specification. Thus, NetApp's construction of this term impermissibly introduces additional limitations not contemplated by the inventors. *Phillips*, 415 F.3d at 1316 ("[O]ur cases recognize that the specification may reveal a special definition given to a claim term by the patentee that differs from the meaning it would otherwise possess. In such cases, the inventor's lexicography governs.").

NetApp's addition of the italicized language is wrong for two additional reasons. First, a number of claims, such as claim 3, already require a communication be received from or sent to a medium access control module, making that additional limitation in NetApp's construction unnecessary. '630 patent, col. 15:12-14. Other claims, such as claim 15, do not include such a limitation. Thus, as to these claims, NetApp is improperly reading a limitation into the claims. *CAE Screenplates*, 224 F.3d at 1317; *Comark Communications*, 156 F.3d at 1187.

b. NetApp's Construction Conflates The Different Uses Of "Elements" In The Claims.

Far from achieving NetApp's stated goal of assisting the jury, NetApp's construction sows confusion on multiple levels. The '630 patent has claims that cover the *transmission* of data and other claims that cover the *receipt* of data. Ennis Decl., ¶ 38. An example of a transmission claim is claim 3. Claim 3 recites "receiving a communication at a distribution module of a network interface device from a medium access control module." '630 patent, col. 15:12-14. Claim 3 then requires that "elements of said communication [are distributed] into multiple portions." *Id.*, col. 15:18-20. A first of these portions of a communication is then sent on a first channel and a second portion is sent on a second channel. *Id*., col. 15:20-28. The claim language thus establishes a sequence of events where elements of a received communication are distributed and, once distributed, become a part of portions of communications that are sent out over a network. Ennis Decl., ¶ 38. NetApp's brief acknowledges this. NetApp Op. Br., p. 33.

This sequence mandates that "elements of a communication" must *preexist* "portions of a communication" because "portions" are formed out of the distributed "elements." '630 patent, col. 15:12-24. Ennis Decl., ¶ 38. Thus, in claims that address transmission of information (like claim 3), "elements of a communication" are initially not a part of any portion of a communication because the portions have not yet been formed. *Id*. Instead, such elements of a communication are simply elements of the data stream that in the future may form a portion of that data stream, as the claim requires. *Id*. NetApp's construction, however, requires the elements of a communication to *already be* a part of a portion of a communication (which NetApp incorrectly construes as a mini-frame) even before that portion comes into being. That, of course, is impossible.

The structure of claims covering the "receipt" of data over a network, such as claim 15, further highlights the error of lumping together the different uses of "element" into one construction. Claim 15 teaches a sequence that generally is the reverse of claim 3. Ennis Decl., ¶ 39. It recites receiving a "first portion of a communication" and a "second portion of a communication" over a first and second channel, respectively. '630 patent, col. 16:41-46. "[A]n

element of said first portion and an element of said second portion" are then collected. *Id*., col. 16:47-48. In *this* context the claimed "elements" are indeed a part of portions of communications that are received over a network. Ennis Decl., ¶ 39. This is because the portions already existed before they were received and before the elements are collected. *Id*.

Thus, NetApp's proposed construction conflates at least two different uses of "elements" in the claims. Where the claims recite "elements of a communication," such elements are not yet a part of any portion of a communication and, therefore, cannot possibly be "a portion of a mini-frame," as required by NetApp. On the other hand, "element of a portion," as indicated by the plain language, refers to elements that are constituents of already-formed portions. Ennis Decl., ¶ 39. Therefore, by imposing a single construction on two different uses of the easily understood word "element," NetApp's creates both error and confusion.

c. NetApp Improperly Requires Every Element Of A Communication To Be Encoded For Transmission.

NetApp improperly adds another limitation requiring every element of a communication be "individually encoded for transmission." Again, nothing in the plain meaning of "element [of a] communication" refers to the element being encoded. Because it is improper to "read into a claim an element which is not present," this proposed extra limitation must be rejected. *Renishaw*, 158 F.3d at 1249.

Moreover, this construction improperly imports limitations from dependent claims into independent claims. For example, claim 3 refers to "elements of [a] communication," but makes no reference to encoding. Encoding is instead recited in claim 5, which depends from claim 3. '630 patent, col. 15:32-35 ("wherein said first physical coding module is configured to encode said apportionment of communication elements"). As established above, importing a limitation from a dependent claim into an independent claim is improper. *Phillips*, 415 F.3d at 1315.

Furthermore, NetApp offers no support for this added requirement other than pointing to the various parts of the specification that discuss encoding of communications. NetApp Op. Br., p. 35. But merely because the specification discusses encoding of communications does not mean that every reference to communications in the claims must encompass that limitation,

especially where, as here, encoding is expressly claimed in only certain dependent claims. Accordingly, NetApp's construction should be rejected.

VI. CONCLUSION

For the above reasons, Sun requests the Court adopt its proposed constructions.

Dated: July 21, 2008

DLA PIPER US LLP

By /s/ Mark Fowler
Mark D. Fowler

Attorneys for Defendant and Counterclaim Plaintiff
SUN MICROSYSTEMS, INC.